ABOVE THE GROUND PLANE

WWVB Preamplifier

COLUMNS

PHOTO 1
Connecting the ends of a 5 meter ribbon cable, offset by one wire, produces a center-tapped 40 turn loop antenna. An epoxy strip molded between Kapton tape reinforces the joint, with strips of transformer tape capturing the taps.

Building a radio means extracting RF from the air, which may mean staying up all night. Ed presents the analog parts of a WWVB receiver.

By Ed Nisley

Radio station WWVB has an entirely predictable playlist: just the time, all the time, without commercial interruptions. Unlike the voice announcements on WWV, its National Institute of Standards and Technology sibling, WWVB transmits time information in machine-readable format, send at 1 baud using pulse-width modulation to encode the 0, 1, and sync marker values. Inexpensive radio-controlled "atomic clocks" synchronize with the WWVB signal and display its BCD digits and status information.

The Arduino-based Totally Featureless Clock I described in my 2010 columns has been hanging on a friend's kitchen wall, quietly decoding the WWVB signal and showing the correct time (and *only* the correct time) in huge blue digits, for the last six years. The C-Max receiver I used in that project has been discontinued, the WWVB signal now includes additional data sent with phase-shift modulation, and I've often thought that building a receiver using GNU Radio would make an interesting project.

The processing inside software-defined radios starts with RF energy from an antenna, so, in this column, I'll describe the analog parts of a WWVB receiver that I'm working on: the antenna and preamplifier that extract a useful signal level from the air.

RESONATING A LOOP

The WWVB transmitters operate at 60 kHz, an octave above the border between the Low Frequency and Very Low Frequency bands, producing a vertically polarized signal with a free-space wavelength of 5 km = 3.1 mi, with the strongest signals occurring during the night. The huge wavelength precludes the simple antenna designs used for radios operating at higher frequencies: a quarter-wave vertical antenna would tower 1.25 km above my house, with a ground plane covering the southern half of the Town of Poughkeepsie.

Signals in the LF and VLF bands suffer considerable interference from man-made radio sources, including the ubiquitous

switching power supplies and inverters found in digital devices. Although EMI filters can suppress the harmonics produced by steep, high-current transitions, stray RF energy from a nearby LED desk lamp can completely wipe out the WWVB signal from Colorado. Worse, cheap devices may not comply with UL and FCC regulations; I've seen PCBs with empty spots where the necessary EMI-filtering components simply weren't installed.

Most WWVB receivers use magnetic loop antennas wound on a ferrite bar for improved local noise rejection in EMI-rich environments. The ferrite core concentrates the RF signal's magnetic field within the surrounding winding, with the effect that a physically small winding can have a much larger effective area, roughly in proportion to the ferrite's relative permeability. Because the ferrite materials useful for LF antennas have relative permeabilities around 500, a 20 mm diameter winding expands to nearly half a meter.

Although I've seen loop antennas wound around bundles of ferrite bars, the antenna's weight increases directly with the number of bars and, if my experience is any guide, its fragility rises exponentially. Having dropped my share of bars and toroids over the years, I know you can't build anything useful from shattered ferrite!

A recent ARRL *QEX* article by John Magliacane, KD2BD, described a WWVB-locked frequency standard that used a simple technique to build a large air-core loop antenna from a length of ordinary ribbon cable: solder the ends together with a one-wire offset, as shown in **Photo 1**. A 40-conductor ribbon cable thus becomes 40-turn loop antenna: 39 solder joints convert a 5 m length of ribbon cable into a 1.6 m diameter loop antenna. Although you must still support the antenna in a roughly circular shape, starting with a ribbon cable eliminates the need to wind 200 m of fine wire into 40 carefully laid turns around a bulky form.

I soldered the ends, reinforced the solder joints by embedding them in quick-curing epoxy, and bound the loose wire ends with strips of transformer tape. Terminating the cable ends with insulation-displacement connectors and joining them through pin headers on a small circuit board would certainly produce a neater and, perhaps, more durable assembly; that's an exercise for the motivated reader.

The windings of a loop antenna form an inductor with stray capacitance between its conductors, which means it will act as a parallel-resonant tank circuit with a resonance peak at the frequency where its inductive and capacitive reactances become numerically equal:

$$f_r = \frac{1}{2\pi\sqrt{L \cdot C_s}}$$

Adding a lumped capacitor at the loop terminals will reduce the resonant frequency and, by matching that frequency to the signal of interest, the antenna response will act as a bandpass filter centered on the loop's resonant frequency.

For a given winding length and number of turns, the loop inductance depends on the shape of the antenna's perimeter, with a circle producing the highest inductance. I measured 6.7 mH for a rough circle and 5.5 mH for a lumpy rectangle, both laid out on the concrete basement floor and measured with an LCR meter.

I suspended the loop from screws driven in the overhead floor joists to produce a droopy rectangle close enough to my workbench that I could connect some test equipment. Driving it from a signal generator through a 470 kΩ isolating resistor and observing the voltage across the loop with an oscilloscope showed a 213 kHz self-resonant frequency.

Knowing the loop's inductance L and self-resonant frequency f_r, its stray capacitance C_s will be:

$$C_s = \frac{1}{\left(2\pi \cdot 213 \times 10^3\right)^2 \cdot \left(6.7 \times 10^{-3}\right)} = 85 \text{ pF}$$

You can cross-check that result by adding capacitance until the resonant frequency drops by a factor of two. Because the frequency varies inversely with the square root of the capacitance, the total capacitance will then be four times the stray capacitance and the stray capacitance equals 1/3 of the added capacitance.

Putting 250 pF across the loop resonated it at 108 kHz, which was close enough to half of 213 kHz for my purposes:

$$83 \text{ pF} = \frac{250 \text{ pF}}{3}$$

Having verified the stray capacitance, the inductance of the loop suspended in mid-air worked out to:

$$L = \frac{1}{\left(2\pi \cdot 213 \times 10^3\right)^2 \cdot \left(83 \times 10^{-12}\right)} = 6.9 \text{ mH}$$

Given the imprecision of my layouts and the measurement tolerances, those numbers hang together quite well!

In order to resonate the antenna at 60 kHz, the total capacitance C_t required will be:

$$C_t = \frac{1}{\left(2\pi \cdot 60 \times 10^3\right)^2 \cdot \left(6.9 \times 10^{-3}\right)} = 1047 \text{ pF}$$

Tuning the loop should therefore require adding a 1047 – 83 = 960 pF external capacitor, so I paralleled capacitors until the resonance hit 59.9 kHz. The 800 pF of added capacitance was lower than I expected, although the total capacitance differs by 10% and, perhaps, I nudged the loop into a different shape over the course of all the manipulations.

The ribbon cable came from eBay and, as is typical of bottom-dollar products, had conductors made from 31 AWG copper of dubious purity, rather than the 28 AWG wires in high-grade cables. Because a 31 AWG conductor has half the area of a 28 AWG conductor, each wire has half the expected copper: less cost for the manufacturer and, most likely, the customer will never notice.

After I soldered the cable ends and before I applied the epoxy, I measured the total resistance to verify end-to-end continuity through all 40 turns and 39 joints. I thought the 80 Ω resistance was surprisingly high, but each 20 turn half-winding measured 40 Ω to the center tap: my solder joints were good and a quick check of the wire tables showed 31 AWG copper actually has that much resistance. Although the cable will work as an antenna, don't get your components from eBay when when you're building something critical.

A tank circuit's Q is the ratio of its reactance at resonance to its resistance:

$$Q = \frac{X_C}{R} = \frac{X_L}{R}$$

The loop antenna's inductive reactance at its 60 kHz resonance is:

$$X_L = 2\pi \cdot f_r \cdot L$$

$$X_L = 2\pi \cdot 60 \times 10^3 \cdot 6.9 \text{ mH} = 2.6 \text{ k}\Omega$$

which means the tank's Q is:

$$Q = \frac{2600}{80} = 32$$

The value of Q also determines the –3 dB bandwidth around the resonant frequency:

$$Q = \frac{f_r}{BW}$$

$$BW = \frac{f_r}{Q} = \frac{60 \times 10^3}{32} = 1.9 \text{ kHz}$$

A better-quality cable with 28 AWG conductors would have half the resistance, twice the Q, and half the bandwidth.

Homework: verify that.

At resonance, the voltage induced in a loop antenna by an RF signal with field strength e in V/m will be:

$$E = Q\frac{2\pi f_r eNA}{c}$$

At the e = 100 µV/m field strength that NIST says should suffice for most WWVB receivers, that works out to 300 µV for my antenna arranged in a vertical circle parallel to the south wall of my basement lab. Obviously, that's insufficient to observe on an oscilloscope without more amplification and filtering, but you can inject a signal from a signal generator driving a small "tickler" loop for testing.

Using the signal generator and oscilloscope, I measured the –3 dB points at 58.1 and 60.1 kHz, giving a 2 kHz bandwidth very close to the expected 1.9 kHz.

Unfortunately, my spectrum analyzer's tracking generator has a 100 kHz lower frequency limit, so I can't present a stimulus-response graph. However, **Photo 2** shows the background noise near my workbench, with the spectrum analyzer's 50 Ω input connected to the loop antenna through a 470 Ω series resistor to reduce its loading effect. The broad hump from about 56 to 62 kHz shows the antenna has increased response in that area, even if the noise doesn't allow for better numbers. The marker sits on a peak at 57.43 kHz from an unknown source that vanished after a few hours.

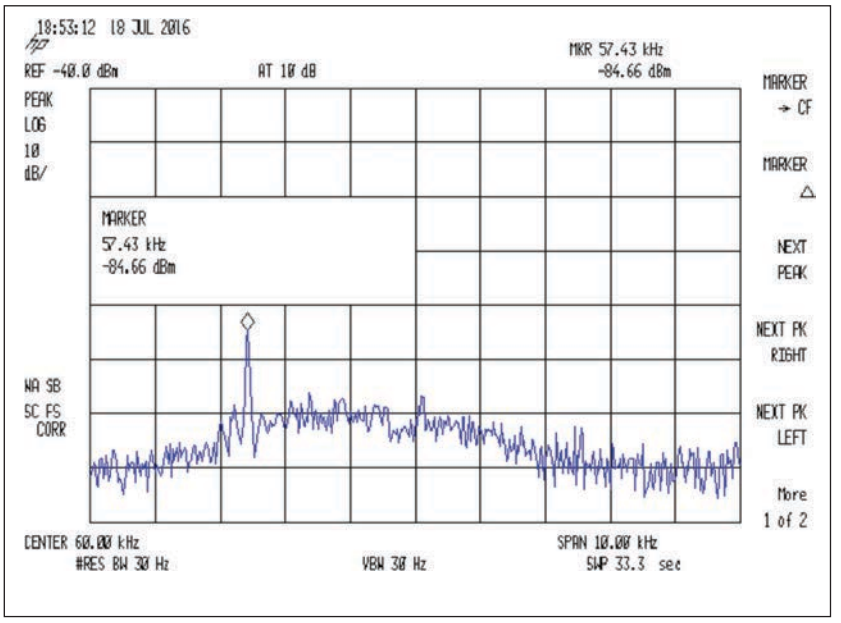

PHOTO 2

Putting about 800 pF across the loop antenna terminals produces a resonant response around 59 kHz, as shown by the noise level in my basement. The 57 kHz spike came from an unknown source and vanished a few hours later.

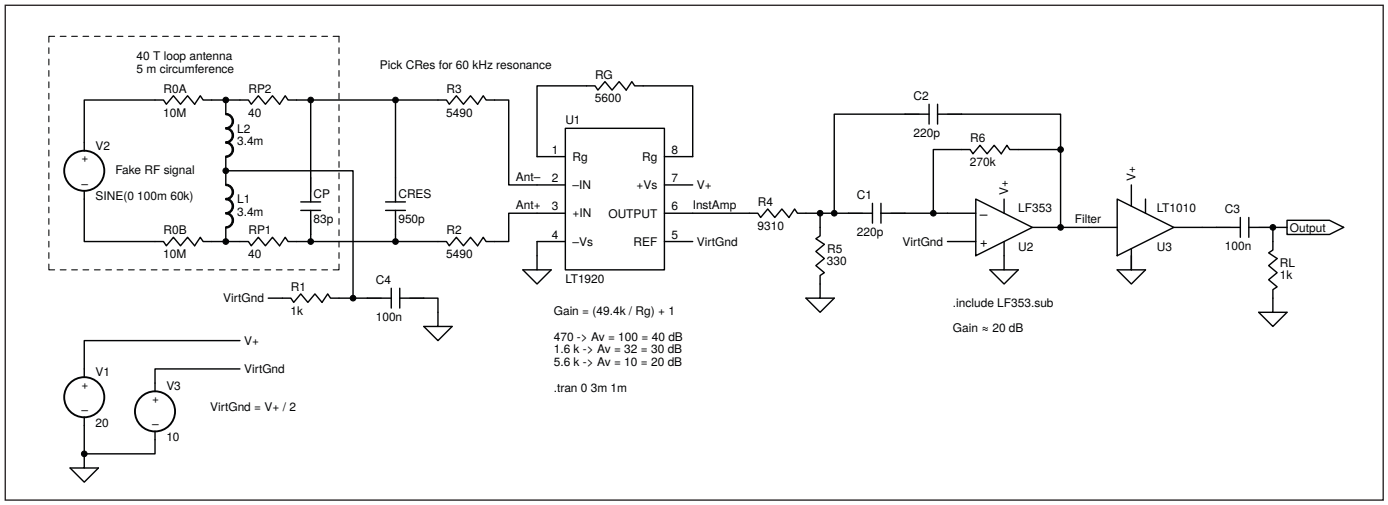

All in all, the loop antenna behaves as it should.

INSTRUMENTATION AMPLIFIER VS LF RF

The desired output of a loop antenna appears as a differential-mode voltage across its terminals that's proportional to the total magnetic field strength parallel to the loop axis within its perimeter. The loop also acts as a simple non-resonant wire antenna responding to the electric field around it by producing a voltage on both conductors: the common-mode voltage of the entire antenna winding varies with respect to the local ground potential. Because nearby EMI sources have disproportionately large electrostatic fields, the RF amplifier must reject the common-mode voltage, with a large ratio between its differential-mode and common-mode gains improving its noise immunity.

The KD2BD design used an RF preamplifier circuit based on an instrumentation amplifier that took advantage of its very high CMRR (Common-Mode Rejection Ratio) to suppress local noise, followed by an op-amp bandpass filter improving the overall selectivity for the differential-mode signals. I substituted present-day parts with somewhat higher performance into that relatively old circuit to produce the Spice schematic in **Figure 1**, giving me a better understanding of how the various sections interacted.

The components inside the dotted box in the upper left simulate the values of the loop antenna, although I didn't match the mutual inductance between the two halves of the loop. The signal source creates a 100 mV_{peak} = 70 mV_{RMS} sine wave that injects a 3.5 µA current into the "antenna". The simulation showed a 250 µV differential voltage across CRes, the added resonating capacitor, roughly equal to the expected voltage from the WWVB signal under nighttime conditions.

A single resistor sets the LT1920 gain, but, unlike most op amps, its frequency response also depends on that resistor: it doesn't have a unique Gain-Bandwidth-Product value. The datasheet shows the gain remains flat to 60 kHz with a voltage gain of 100 = 40 dB, which should be more than enough, and extends to 500 kHz at lower gains.

The datasheet recommends 5 kΩ resistors at the LT1920 inputs to prevent damage from electrostatic discharge. Because lightning strikes can induce huge transient voltages in antennas, I decided that more protection was better; the original KD2BD design used opposed Zener diodes that you'll see in the actual hardware.

The resistors and capacitors surrounding the LF353 op amp implement a multiple-feedback bandpass filter centered on 60 kHz. The original design used a different op amp, so I adjusted the component values to match

FIGURE 1
The LT1920 instrumentation amplifier converts the differential RF signal from the (simulated) resonant magnetic loop antenna into a single-ended signal around the virtual ground at V+/2. The components around the LM353 op amp implement a simple 60 kHz bandpass filter and the LT1010 buffer drives the output cable. Most of the component values match the as-built circuit.

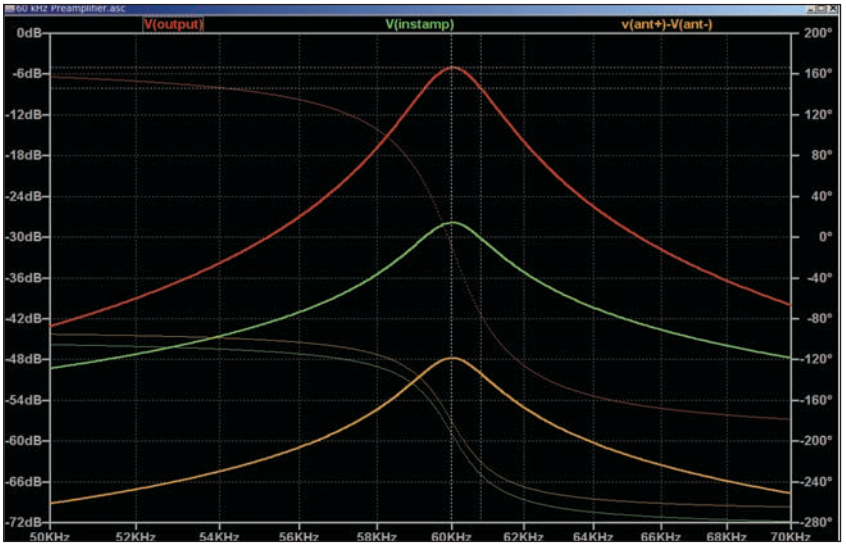

PHOTO 3
The LT1920 and LM353 each provide about 20 dB of gain at 60 kHz, with the resonant antenna and bandpass filter limiting the response to a 1.5 kHz bandwidth. The 20 kHz plot span exaggerates the filter response.

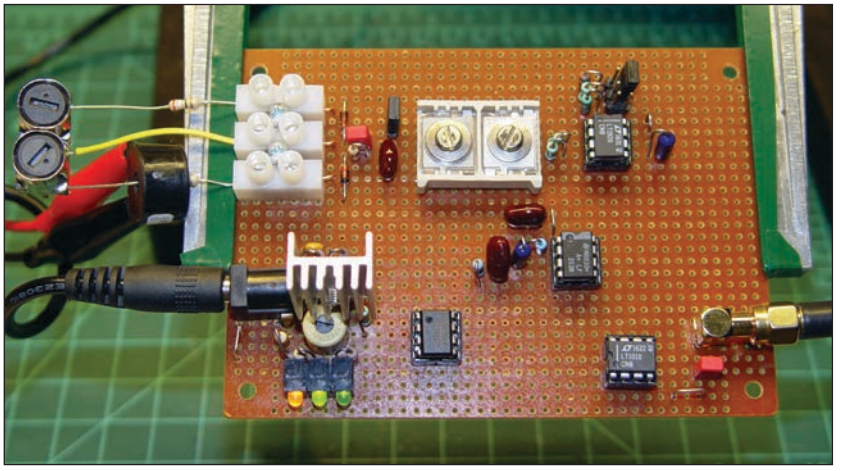

PHOTO 4
The bulky white mica compression capacitors resonate the loop antenna at 60 kHz, with the bandpass filter just below them providing additional selectivity. The inductors and resistors on the antenna input approximate the distributed values of the actual loop antenna, with a current transformer injecting a signal input.

the capacitors I had on hand and tweaked the resistors for 20 dB of gain at 60 kHz. The bandwidth works out to about 2 kHz: Q = 30.

The LT1010 buffer amplifier drives the cable capacitance between the SMA connector and the SDR radio input. That cable may be quite long, as the connection between the preamp and the antenna must use relatively short wires, and the best location for the antenna won't be in the most convenient location for the receiver.

The LT1920 specs call for high power supply voltages, compared to the single +5 V supply found in many almost-all-digital devices. The datasheet also shows that the output swing drops from 28 VPP in the audio band to only 5 V_{PP} = 2.5 V_{peak} at 60 kHz with ±15 V supplies, so lower supplies would exacerbate that limitation.

I ran the amplifiers from a +20 V power supply, with a "virtual ground" at +10 V, mostly because my parts heap produced a 24 VDC wall wart. The center tap of the antenna sits at AC ground, with a 1 kΩ resistor providing a DC path for the LT1920's input bias currents to the virtual ground.

The simulated response in **Photo 3** shows a 1.5 kHz bandwidth from the antenna input (yellow trace) to the output (red trace). The LT1920 (green trace) has a flat response in this range, so it's a uniform 20 dB above the input. The bandpass filter provides slightly over 20 dB gain at 60 kHz and rolls off to either side.

All of that seemed reasonable, so the next step was to actually build something.

HARDWARE CONSTRUCTION

I hand-wired the circuit on the prototyping board in **Photo 4**, leaving plenty of room for modifications and additions. I used DIP versions of the op amps, because they're easier to hand-wire, and put them in sockets, because semiconductors deliberately exposed to high transient voltages may need replacement.

The KD2BD preamp ran on +12 V supplied from the main receiver through the coaxial cable, with homebrew bias tees separating the RF signal and DC voltage. I opted for a local power supply, because adding a bias tee to an SDR radio seemed overly complex.

The power supply components clustered around the heatsink in the lower left corner of Photo 4 implement the two voltage sources in the simulation. The schematic in **Figure 2** shows an LM317 regulating the 24 V from the wall wart down to 20 V and an LF411 producing the low-impedance virtual ground at half the LM317's output. The three LEDs along the lower edge of the board indicate that all three supply voltages are functional and, because the LM317 requires a minimum output current to ensure proper regulation, D4 serves as both an indicator and a small load.

The preamplifier circuit in **Figure 3** closely follows the simulated schematic. I used silver-mica capacitors in the RF path, along with an antique dual-stage mica compression capacitor to resonate the loop, although contemporary film capacitors and trimmers should work equally well.

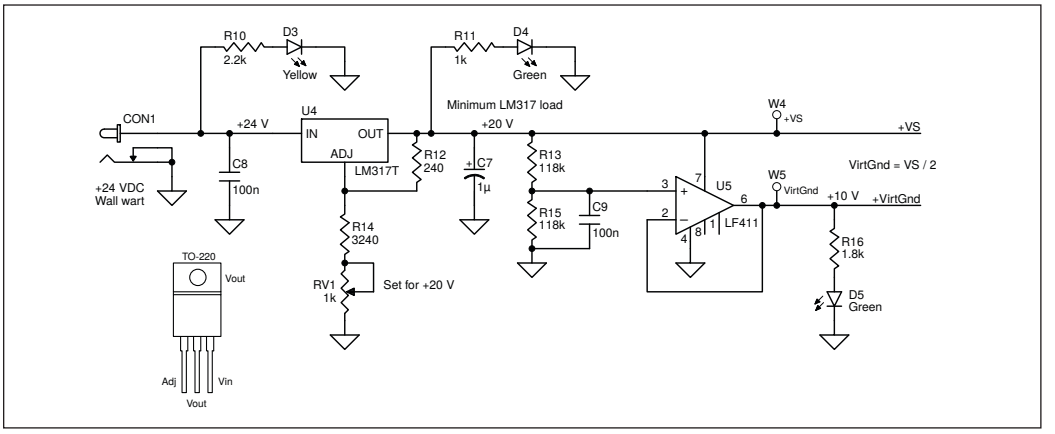

FIGURE 2
The power supply centers a virtual ground at +10 V, half of the 20 V from the LM317 , so the rest of the circuitry thinks it's operating from a ±10 V split supply. D4 provides a minimum load to keep LM317 happy, even with all of the other ICs unplugged from their sockets. The LM411 op amp provides bias currents to the other amplifiers to set their outputs at +10 V.

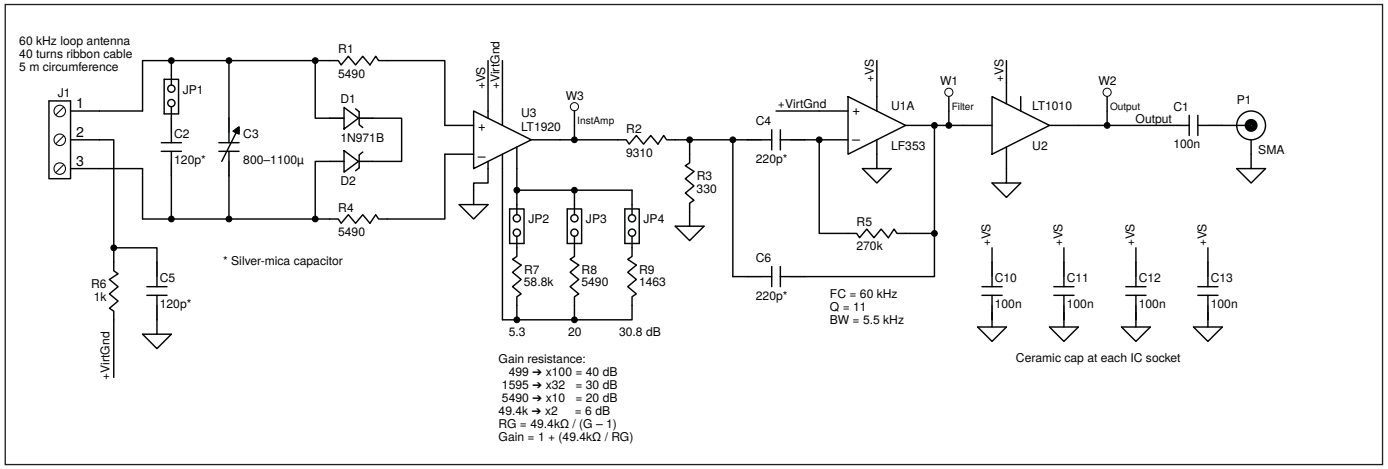

FIGURE 3
The three jumpers below the LT1920 set its gain appropriate to the minimum signal strength. The components came from my parts bin, with their values as measured.

The three resistors below the LT1920 set its gain to selected values between 0 dB (no jumper) and 30 dB. Because I don't have a complete E192 set of resistors, I used the closest match to the ideal values; the varying RF input means an exact gain isn't critical.

Some early testing showed that 40 dB of gain turned the LT1920 into an oscillator. That wasn't entirely surprising, due to the total lack of shielding and likelihood of feedback through the virtual ground. I proceeded with 20 dB of gain for most of the testing.

The adjustable inductors and resistors connected to the antenna screw terminals in the upper left of Photo 4 correspond to the antenna components in the simulation. The small current transformer around the lower resistor injects a differential current from a signal generator into the circuit.

That test setup produced the interesting result in **Photo 5**. The central peak comes from a DDS signal generator set to 59.975 kHz, with the spectrum analyzer's 30 Hz resolution bandwidth filter broadening the single-frequency input. I didn't expect the two peaks surrounding it, however, due to the 60 Hz ambient power-line magnetic field modulating the "carrier" signal. The sidebands occur at 120 Hz offsets, because the magnetic field strength reaches a maximum during each half of the 60 Hz power line current. Wider frequency spans use broader RBW filter settings that obscure those details.

I adjusted the compression capacitors for peak response with the fake antenna, then connected the loop antenna to the preamp. A few turns of wire wrapped around the ribbon cable injected a 60.000 kHz test signal from the DDS signal generator and, with the capacitors already close to the proper setting, I tuned for maximum response from the actual antenna.

Measuring the end-to-end frequency response required a trick that produced **Photo 6**, which isn't quite the sequence of resonances it appears to be. Remember that the analyzer's tracking generator doesn't cover the preamp's frequency range.

I set the analyzer to display-and-hold the maximum response at each frequency, then adjusted the DDS generator's output frequency in 250 Hz steps from 57.5 kHz to 62.5 kHz. Each of the peaks in Photo 6 corresponds to the preamp response at one of those input signals, with the spectrum analyzer displaying all of them in one trace. Smaller frequency steps would produce more peaks located closer together, with the ideal being the smooth curve skating along their top from a tracking generator.

Even with 250 Hz steps, however, you can see the preamp's overall response. The left marker sits 1 kHz below the filter's 60 kHz center at –3 dB. The measured 2 kHz bandwidth corresponds reasonably well with the 1.5 kHz bandwidth from the simulation,

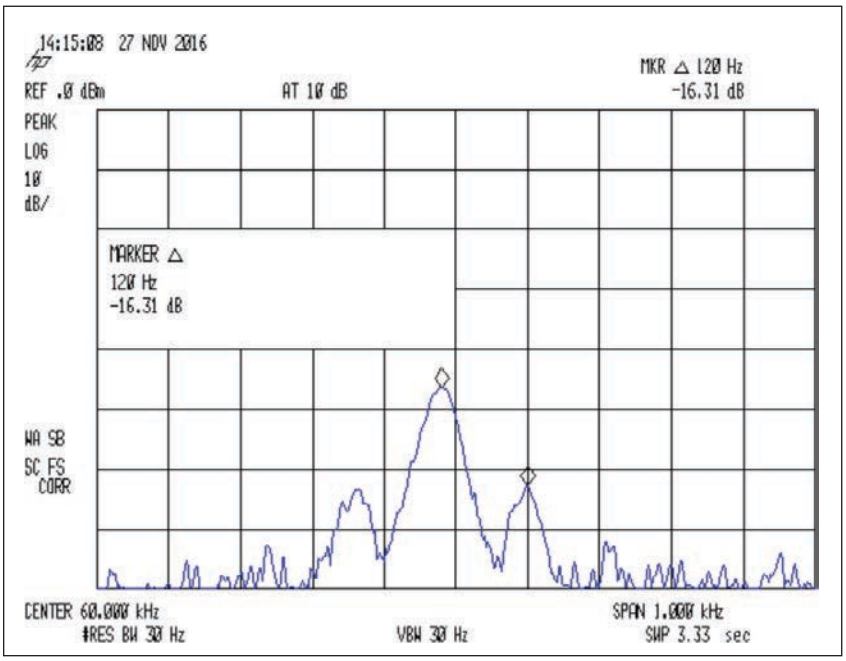

PHOTO 5
The sidebands around the central peak show the ambient magnetic field modulating a test signal near 60 kHz.

COLUMNS

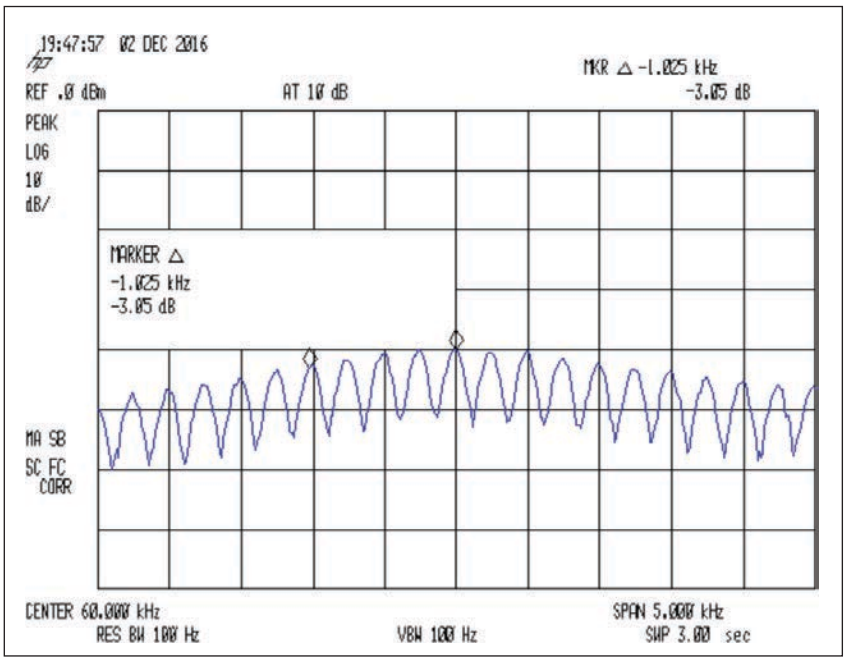

PHOTO 6
This image overlays 21 separate measurements of the preamp's response to an input signal, with successive test frequencies spaced at 250 Hz. The maximum value of each peak represents the response at that frequency, with the left marker showing the -3 dB point of the response at 59 kHz.

both of which depend on the characteristics of their respective antennas.

The various gain jumper settings produced the expected results, although the 30 dB setting occasionally pushed the preamp into oscillation. Because the preamp definitely oscillates with 40 dB of gain using the fake antenna, I think the loop antenna picks up feedback from the preamp's rather casual layout. The original KD2BD design used an instrumentation amplifier with 20 dB of gain and no shielding, so perhaps my circuit behaves similarly.

FAILURE TO RECEIVE

During the Totally Featureless Clock project, I wired a C-Max WWVB receiver to a single red LED, powered from a pair of alkaline D cells, as a failsafe propagation monitor: it's a pure analog receiver that requires only a radio signal. I've gotten used to watching its LED blink calmly at 1 Hz, at least during the morning and evening hours, and it's easy to distinguish the time code's 0, 1, and sync pulses, even if I can't accumulate all the BCD digits in my head. Erratic flickering indicates the usual loss of propagation during daylight hours, when there's insufficient RF for good reception.

Based on that, I've discovered any attempt to build a WWVB project on my workbench can destroy LF radio propagation from Colorado. During the course of building this preamp in November, I seem to have triggered a major early winter storm over the Midwest, covered upstate New York with several feet of snow, and reduced the WWVB signal to the point where my failsafe receiver flickered feebly during the morning and evening hours when I normally see good reception.

Radio-synchronized clocks continued to work, because they need only a few minutes of good reception and, deep in the night, I'm sure that happened on a regular basis. Intermittent reception isn't nearly enough for my purposes and, being unwilling to stay up all night for science, I haven't been able to verify the entire signal path from the transmitters in Colorado, through the antenna and preamp, to the spectrum analyzer on my workbench.

The next step will be to move the antenna and preamp to the attic, far from the local EMI sources, while awaiting better propagation. A Raspberry Pi 3 in the attic provides a receive-only APRS iGate through a VHF radio, so I should be able to install GNU Radio and work on a remotely accessible WWVB receiver.

CONTACT RELEASE

The files for this column include the LTSpice simulation model and the Kicad schematic; for obvious reasons, there is no PCB layout.

ABOUT THE AUTHOR

Ed Nisley is an EE and author in Poughkeepsie, NY. Contact him at ed.nisley@pobox.com with "Circuit Cellar" in the subject line to avoid spam filters.

circuitcellar.com/ccmaterials

RESOURCES

Ground Wave Propagation, https://en.wikipedia.org/wiki/Ground_wave_propagation.

Loop antenna theory, http://sidstation.loudet.org/antenna-theory-en.xhtml.

J. A. Magliacane (KD2BD), "A Frequency Standard for Today's WWVB," *QEX*, November/December, 2015, www.arrl.org/qex.

E. Nisley, "Totally Featureless Clock," Parts 1–3, *Circuit Cellar* 235/237/239, 2010.

———, "Can't Set Your Clock," *Circuit Cellar* 131, 2001.

NIST, Radio Station WWVB, www.nist.gov/time-and-frequency-services/nist-radio-stations/wwvb.

SOURCE

LT1920 instrumentation amplifier
Linear Technology | www.linear.com